

CLAIMS

What is claimed is:

- 5      1. A method of accessing memory from multiple interconnected access nodes, at least one access node including a memory unit accessible by other interconnected access nodes, the method comprising:
  - receiving a memory access request at a given access node from a source other than from another access node;
  - 10        determining a physical memory address associated with the memory access request; and
  - determining whether the physical memory address corresponds to a memory unit local to the given access node and:
    - if so, servicing the memory access request by accessing a memory unit local to the given access node;
    - 15        if not, servicing the memory access request by accessing a memory unit located at a node other than the given access node.
2. A method as in claim 1, wherein determining a physical memory address includes:
  - at the given node receiving the access request, converting a logical address associated with the access request to the physical address.
3. A method as in claim 1 further comprising:
  - 25        at each of multiple interconnected access nodes, maintaining an allocation table to map logical addresses to corresponding physical addresses identifying to which of the multiple interconnected access nodes the logical address pertains.
4. A method as in claim 1, wherein accessing a memory unit located at a node other than the given access node includes:
  - 30

via a serial link, supporting communications between the given access node and a memory unit at another access node over a physical point-to-point connection.

5 5. A method as in claim 1 further comprising:

    fabricating the access nodes as individual circuit boards that plug into a common backplane, the common backplane including physical connections to link each of the access nodes to each other.

10 6. A method as in claim 1 further comprising:

    storing portions of a logical data stream in memory units of multiple interconnected access nodes.

7. A method as in claim 6 further comprising:

15      at the given access node, retrieving selected portions of the logical data stream over time; and

    transmitting the logical data stream to a user.

8. A method as in claim 7 further comprising:

20      receiving the logical data stream for playback in real-time.

9. A method as in claim 6 further comprising:

    periodically generating memory access requests to retrieve portions of the logical data stream.

25

10. A method as in claim 1 further comprising:

    at the given access node, receiving local memory access requests from the source;

30      at the given access node, receiving remote memory access requests from other access nodes; and

scheduling times for servicing the local memory access requests and remote memory access requests.

11. A method as in claim 1, wherein receiving a memory access request includes:
  - 5 receiving a request to read from the local memory unit of the given access node;
  - receiving a request to write to the local memory access unit of the given access node; and
  - scheduling times for servicing the requests to read and write from/to the local memory unit of the given access node.
12. A data storage system comprising:
  - a local access node including:
    - 15 a node interface to support communications with multiple remote access nodes and their corresponding memory units;
    - a local memory unit to store data;
    - a port interface to receive memory access requests from a source other than the remote access nodes; and
    - 20 a translator to process the memory access requests received through the port interface, the translator utilizing a data identifier associated with a given memory access request to determine whether to access the local memory unit or one of the multiple remote access nodes' memory units in order to service the given memory access request.
  - 25 13. A data storage system as in claim 12, wherein the translator converts the data identifier associated with the given memory access request to a physical address identifying, at least in part, which memory unit of the corresponding multiple remote access nodes to access in order to service the given memory access request.

-39-

14. A data storage system as in claim 12 further comprising:  
links over which the remote access nodes and the local access node  
communicate to access data stored in their corresponding memory units.
  
- 5 15. A data storage system as in claim 14, wherein the links include point-to-point  
connections physically coupling each remote and local access nodes to each other.
  
- 10 16. A data storage system as in claim 14, wherein parallel data retrieved from a  
memory unit of a remote access node is converted to a serial stream prior to  
transmission over a link to the local access node.
  
- 15 17. A data storage system as in claim 14, wherein the access nodes are circuit boards  
that plug into a common backplane, the common backplane including physical  
connections to link the access nodes to each other.
  
18. A data storage system as in claim 12, wherein the memory units of the access  
nodes include electronically addressable memory chips.
  
19. A data storage system as in claim 12, wherein the data identifier associated with a  
20 given memory access request includes a request for a logical data stream stored, at  
least in part, among multiple remote access nodes.
  
20. A data storage system as in claim 12, wherein the data identifier associated with a  
given memory access request includes a request for a logical data stream, portions  
25 of which are stored in memory units of both the local and remote access nodes.
  
21. A data storage system as in claim 20, wherein portions of the logical data stream  
are timely retrieved from the remote and local access nodes for substantial real-  
time playback of the logical data stream.

-40-

22. A data storage system as in claim 21, wherein the port interface periodically generates memory access requests to retrieve portions of the logical data stream that are, in turn, transmitted to a user substantially playing back the logical data stream in real-time.

5

23. A data storage system as in claim 22, wherein the local access node further includes:  
an arbitrator to arbitrate servicing of locally received memory access requests received through the port interface of the local access node and remotely received memory access requests received through the node interface from remote access nodes.

10

24. A data storage device as in claim 23, wherein the arbitrator includes:  
a scheduler to schedule locally and remotely received memory access requests for access to the local memory unit.

15

25. A data storage system as in claim 12, wherein the memory access requests include both requests to read and requests to write to corresponding memory units of the access nodes.

20

26. A data storage system as in claim 12, wherein contents of the local memory unit and remote access nodes' memory units are maintained by both the local and remote access nodes to collectively form a shared memory storage system.

25

27. A data storage system as in claim 26, wherein the translator includes an allocation table for mapping logical addresses associated with stored portions of a data stream to physical addresses identifying a corresponding location of the shared memory storage system in which to locate portions of the data stream.

28. In a data storage system including multiple access nodes and corresponding memory units interconnected via links, an access node comprising:

- a memory unit to store data information;
- an interconnect interface circuit supporting communications with the multiple remote access nodes and their corresponding memory units over the links;
- a port interface to receive memory access requests from a source other than remote access nodes communicating through the interconnect interface circuit; and
- 10 a translator circuit that processes a given memory access request received from the source through the port interface to determine which of the multiple access node's memory unit of the data storage system to access in order to service the given memory access request.

15 29. A data storage system comprising:

- a local access node including:
  - a node interface to support communications with multiple remote access nodes and their corresponding memory units;
  - a local memory unit to store data;
  - 20 a port interface to receive memory access requests from a source other than the remote access nodes; and
  - means for processing the memory access requests received through the port interface, the processing means translating a data identifier associated with a given memory access request into a physical memory address to determine whether to access the local memory unit or one of the multiple remote access nodes' memory units in order to service the given memory access request.

25 30. A computer program product including a computer-readable medium having instructions stored thereon for processing data information, such that the

instructions, when carried out by a processing device, enable the processing device to perform the steps of:

receiving a memory access request at a given access node from a source other than from another access node;

5 determining a physical memory address associated with the memory access request; and

determining whether the physical memory address points to a memory unit local to the given access node and:

10 if so, servicing the memory access request by accessing a memory unit local to the given access node;

if not, servicing the memory access request by accessing a memory unit located at a node other than the given access node.

31. A data storage system comprising:

15 a plurality of memory access nodes, each of the memory access nodes including an interconnect interface to support communications with other memory access nodes; and

links interconnecting the plurality of memory access nodes,

20 at least one of the plurality of memory access nodes further including a memory unit for storing data,

at least one of the plurality of memory access nodes further including a port interface module for communications with a user application initiating access to stored information,

25 each of the plurality of access nodes further including a memory access arbiter for arbitrating communications to at least one of a local memory unit and a remote memory unit of another access node.

32. A data storage system as in claim 31, wherein the links comprise point-to-point connections physically coupling the memory access nodes to each other.

-43-

33. A data storage system as in claim 31, wherein the memory access nodes comprise circuit boards that plug into a common backplane, the common backplane including physical connections to link the memory access nodes to each other.

5 34. A data storage system as in claim 31, wherein the local memory unit and remote memory unit include electronically addressable memory chips.

10 35. A data storage system as in claim 31, wherein the electronically addressable memory chips are solid- state RAM (Random Access Memory) based memory devices.

15 36. A data storage system as in claim 31, wherein the arbiter further determines which of at least one memory unit to access to service a memory access request received from the user application.

37. A data storage system as in claim 36, wherein the memory access request includes a request to read from and/or a request to write to memory.

20 38. A data storage system as in claim 36, wherein the memory access request includes a data identifier and a logical address.

25 39. A data storage system as in claim 38, wherein the memory arbiter further includes an address translator to convert the data identifier and the logical address into a corresponding at least one physical address identifying which of at least one memory unit to access to service the memory access request.

40. A data storage system as in claim 39, wherein the address translator includes an allocation table for mapping logical addresses to the corresponding at least one physical address.

41. A data storage system as in claim 38, wherein the data identifier includes a request for a logical data stream, a portion of which is stored in the local memory unit and another portion of which is stored in the remote memory unit.

5 42. A data storage system as in claim 41, wherein the portions of the logical data stream are timely retrieved from the local memory unit and the remote memory unit for substantial real-time playback of the logical data stream.

10 43. A data storage system as in claim 42, wherein the port interface periodically generates memory access requests to retrieve portions of the logical data stream that are, in turn, transmitted to the user application substantially playing back the logical data stream in real-time.

15 44. A data storage system as in claim 31, wherein the memory arbiter further schedules times for servicing a plurality of memory access requests from users.

20 45. A data storage system as in claim 31, wherein the interconnect interface further includes a transceiver for processing outgoing and incoming data communicated over the links, the transceiver including a parallel-to-serial transmitter for transmitting data over the links and a serial-to-parallel receiver for receiving data over the links.

25 46. A data storage system as in claim 31, wherein the port interface supports communication via an application-level memory access protocol, and wherein each memory access node further includes an interconnect interface that supports communication using an inter-nodal memory access protocol.

47. A data storage structure comprising:

a plurality of data storage systems according to claim 31, each of the plurality of data storage systems having a memory access node to support communications with other data storage systems; and  
links interconnecting the plurality of data storage systems.

5

48. In a data storage system including a plurality of memory access nodes and links interconnecting the plurality of memory access nodes, at least one memory access node comprising:

a local memory unit for storing data;

10 an interconnect interface for supporting communications with other memory access nodes over the links;

a port interface module supporting communications with a user application initiating access to stored information; and

15 an arbiter for arbitrating communications among the memory unit, the interconnect interface and the port interface module, the arbiter determining whether to access the local memory unit or at least one other memory unit in the other memory access nodes to service a memory access request associated with the user application.

20 49. The at least one memory access node as in claim 48, wherein the port interface module supports communication via an application-level memory access protocol, and wherein the interconnect interface supports communication with the other memory access nodes using an inter-nodal memory access protocol.

25 50. The at least one memory access node as in claim 49, wherein the memory access request includes a data identifier and a logical address, and the arbiter further includes an address translator for converting the data identifier and the logical address into corresponding at least one physical address identifying which of the at least one memory units to access to service the memory access request.

30

51. The at least one memory access node as in claim 50, wherein the address translator includes an allocation table for mapping logical addresses to corresponding physical addresses.
- 5 52. The at least one memory access node as in claim 48, wherein the arbiter further schedules times for a plurality of memory access requests from users.
- 10 53. A data storage system comprising:
  - a plurality of memory access nodes; and
  - links interconnecting the plurality of memory access nodes, each of the plurality of memory access nodes including:
    - a memory means for electronically storing data;
    - a first interface means for supporting communications with the other memory access nodes;
  - 15 a second interface means for supporting communications with a user; and
  - an arbiter means for arbitrating communications among the memory means, the first interface means, and the second interface means, the arbiter means determining which memory means to access to service a memory access request from the user.
- 20 54. A data storage system as in claim 53, wherein the arbiter means translates a logical address associated with the memory access request into a corresponding physical address identifying which memory unit to access to service the memory access request.
- 25 55. A data storage system comprising:
  - a plurality of memory access nodes, each of the memory access nodes including an interconnect interface to support communications with other memory access nodes; and
  - 30 links interconnecting the plurality of memory access nodes,

at least one of the plurality of memory access nodes further including a memory unit for storing data,

5 at least one of the plurality of memory access nodes further including a port interface module for communications with a user application initiating access to stored information,

10 each of the plurality of access nodes further including a memory access arbiter for arbitrating communications to at least one of a local memory unit and a remote memory unit of another access node, the data storage system supporting operations of: receiving a memory access request from a user at a given access node via a port interface module;

determining a memory address associated with the memory access request;

15 determining whether the memory address points to a memory unit local to the given access node; and

20 if the memory address points to a memory unit local to the given access node, servicing the memory access request by accessing a memory unit local to the given access node; and

25 if not, servicing the memory access request by accessing a memory unit located at a node other than the given access node.

56. The data storage system as in claim 55, wherein determining a memory address includes:

25 converting a logical address associated with the memory access request into a physical address identifying which memory unit to access to service the memory access request.

57. The data storage system as in claim 56, wherein converting a logical address into a physical address includes utilizing an allocation table mapping logical addresses to corresponding physical addresses.

-48-

58. The data storage system as in claim 55, wherein accessing a memory unit located at a node other than the given access node further includes:

converting parallel data into corresponding serial data prior to communication over the links.

5

59. A data storage system comprising:

a plurality of memory access nodes, each of the memory access nodes including an interconnect interface to support communications with other memory access nodes; and

10

links interconnecting the plurality of memory access nodes, at least one of the plurality of memory access nodes further including a memory unit for storing data,

at least one of the plurality of memory access nodes further including a port interface module for communications with a user application initiating access

15

to stored information,

each of the plurality of access nodes further including a memory access arbiter for arbitrating communications to at least one of a local memory unit and a remote memory unit of another access node, the data storage system supporting operations of: storing portions of a logical data stream in multiple memory units in the data storage system;

20

receiving a memory access request for the logical data stream from a user at a given access node via a port interface module;

determining memory addresses associated with the memory access request;

25

retrieving the portions of the logical data stream from the multiple memory units over time; and

transmitting the logical data stream to the user at the given access node.

60. The data storage system as in claim 59, wherein transmitting the logical data stream further comprises transmitting the logical data stream for playback in real-time.

5 61. The data storage system as in claim 59 further supporting an operation of:  
periodically generating memory access requests to retrieve the portions of the logical data stream.

62. A data storage system comprising:  
10 a plurality of memory access nodes, each of the memory access nodes including an interconnect interface to support communications with other memory access nodes; and

links interconnecting the plurality of memory access nodes,  
at least one of the plurality of memory access nodes further including a

15 memory unit for storing data,  
at least one of the plurality of memory access nodes further including a port interface module for communications with a user application initiating access to stored information,

each of the plurality of access nodes further including a memory access arbiter for arbitrating communications to at least one of a local memory unit and a 20 remote memory unit of another access node, the data storage system supporting operations of:

at a given access node, receiving local memory access requests via port interface modules;

25 at the given access node, receiving remote memory access requests from other access nodes; and

scheduling times for servicing the local memory access requests and remote memory access requests.

30 63. A data storage system comprising:

a plurality of memory access nodes, each of the memory access nodes including an interconnect interface to support communications with other memory access nodes; and

links interconnecting the plurality of memory access nodes,

5 at least one of the plurality of memory access nodes further including a  
memory unit for storing data,

at least one of the plurality of memory access nodes further including a port interface module for communications with a user application initiating access to stored information,

10 each of the plurality of access nodes further including a memory access arbiter for arbitrating communications to at least one of a local memory unit and a remote memory unit of another access node, the data storage system supporting operations of: receiving a request to read from the local memory unit of the given access node;

15 receiving a request to write to the local memory unit of the given  
access node; and

scheduling times for servicing the requests to read from and write to the local memory unit of the given access node.

20 64. A computer program product including a computer-readable medium having instructions stored thereon for processing data information, such that the instructions, when carried out by a processing device, enable the processing device to perform the steps of:

receiving a memory access request at a given access node from a user;

25 determining a physical memory address associated with the memory access request;

determining whether the physical memory address points to a memory unit local to the given access node; and

30 if so, servicing the memory access request by accessing a memory unit local to the given access node;

-51-

if not, servicing the memory access request by accessing a memory unit located at a node other than the given access node.